

Page 4, delete the first paragraph and insert therefor:

A2
Multicast routing protocols are classified as "sparse mode" or "dense mode". In sparse mode, reception of a multicast transmission by a receiver is accomplished by a multicast join, which is accomplished by a group membership report, which is a message sent from the receiver to the nearest router (the so called "first hop router"), requesting membership in the multicast group, and thus the transmission. If the router is already part of the multicast tree and is already receiving the transmission, then the transmission is simply routed to the new receiver. If not, the router sends a join message to the next router in the chain going to either the source (if known) or a rendezvous point (RP, also called a Core), if the location of the source is not known. The join request travels towards the source or the RP, either a router is reached that is already receiving the multicast transmissions, or until the source or the RP is reached. In sparse mode a receiver stops receiving a multicast transmission (i.e., leaves the multicast group), by sending a "prune" message to the first hop router, which then ceases forwarding transmissions to the receiver. Multicast state in the routers is always subject to timers, and required periodic refreshing to remain valid; because of this it is also possible to stop receiving multicast transmissions by remaining silent, and thus to "time out". In either case, each router in the tree, if it is no longer forwarding the multicast transmission to any receiver, will itself send a prune message to next router in the tree to be removed from the tree entirely.

Page 5, delete the first six lines and insert therefor:

A3
“flood” stage), which then have to explicitly prune themselves if they are not interested in receiving the transmissions (the “prune” stage). This “flood and prune” technique, although technically easy to implement, and is in commercial use on small Internets, or small subsets of large Internets, is not suited for deployment on arbitrarily large Internets due to the geometrical multiplication of the data transmissions required during the initial flood stage.

Page 5, delete the second paragraph, and page 6, delete first three lines, and insert therefor:

A4
Unicast transmissions of data between Autonomous Systems are done through the use of specially chosen routers known as Border Routers, or BRs, with a Border Gateway Protocol (BGP) or similar mechanism to facilitate the exchange of unicast routing information between different Autonomous Systems. Using these protocols, a BR in one Autonomous System can discover whether a node address exists in another Autonomous System and details on unicast routing to that other node address in the other Autonomous System.

Page 7, delete first 7 lines and insert therefor:

A5
there be multiplied into many separate streams. Accounting for this multicast traffic, so that AS1 can properly pay AS2 for the work entailed by this transmission, would thus require detailed knowledge of the internal traffic within AS2, and this might reveal proprietary information about AS2. An audit of this accounting could not be done without intrusive

AS
monitoring of conditions within AS2, which would also reveal sensitive and proprietary information about the workings of AS2 to AS1.

Page 8, delete first nine lines and insert therefor:

AS
also less efficient for situations where there are multiple multicast sources emanating from one IP address, in that there has to be a separate multicast tree maintained for each such source. It is also possible that there would be three or more Autonomous Systems involved in a multicast transmission, say AS1, AS2 and AS3. In this situation, all sources might be located in AS1 and all receivers in AS3, but AS2 might be essential in the construction of the multicast tree between AS1 and AS3. AS2 is thus forced with performing work for which it has no customers, and thus no commercial reason to perform. This problem is called a "third party dependency" in the literature.

Page 9, delete the fifth paragraph and insert therefor:

AS
In yet another particularly advantageous embodiment of the invention, TTP is set up to prevent unauthorized sources from transmitting data on the Trusted Third Party Network (TTPN) through the use of a Multicast Firewall.

Page 11, delete the third paragraph, and page 12, delete the first five lines, and insert therefor:

A⁸
It is generally thought that not switching to a SPT for each source is not good practice, and, indeed, in an Internet the general practice is to set the threshold to zero, so that the transfer to the source based SPT occurs immediately. This is done to avoid having a "hot spot" at the RP, which would have to handle the routing for all sources in the group, and to keep the multicast traffic exclusively on SPTs. In the case of the one-to-many static transmissions on a TTPN, these benefits are illusory, as all traffic will use a SPT, and because some router would have to be the first hop router for the TTPN traffic, and would thus have to accommodate all of the TTPN traffic in any case.

Page 13, delete lines 5 through 12 and insert therefor:

- A⁹
- Rejection of PIM join/leave messages which refer to multicast groups or sources not used by the TTPN.
 - Rejection of Multicast traffic from outside (i.e., all multicast traffic flows outward only).

The use of this multicast firewall will thus protect against unauthorized transmissions from the TTPN facilities (including transmissions from elsewhere that would then otherwise be immediately switched to a SPT not using the TTPN RP), as well as unauthorized termination of the TTPN transmissions.

Page 15, delete the second paragraph and insert therefor:

A10
In that case, the separate streams will have uncorrelated packet losses, and, for a given mean loss rate, ϵ , and N separate staggered streams of equal size, the total loss rate will be ϵ^N . In actual practice, a mean loss rate 10 % (or $\epsilon = 0.1$) represents a high rate of packet loss, and a typical value for packet loss correlation time would be 1 second. Given these parameters, the following table describes the expected performance of a SEC system:

Page 16, delete lines 20-23 and insert therefor:

All
with a typical drop-out duration being 1 second. A SEC with three staggered streams is sufficient to reduce drop outs to the level of a few per hour under fairly extreme conditions of packet loss, and to near zero at other times, which is a reasonable loss rate for many applications.

Page 17, delete the first paragraph, page 18 delete the first line, and insert therefor:

A12
In the most straight-forward SEC implementation, each staggered stream would be a full copy of the original data. The above SEC implementation, with three staggered streams, would reduce drop-outs to a few per hour in the above example, but at the cost of tripling the bandwidth required. In many cases, however, such as most audio and video transmissions for entertainment, a degraded copy of the data stream, at a reduced bandwidth, may be an acceptable replacement for the full data rate. In audio entertainment, for example, using psycho-acoustic compression, while a bandwidth of 160 kilobits per second is required to give full sound quality, a bandwidth of 64 kilobits per second still provides acceptable stereo sound reproduction at most times, and a

bandwidth of 32 kilobits per second is marginally acceptable in monaural sound reproduction.

A12
Since a stereo sound reproduction can be sent as two monaural reproductions, the staggered channels used in SEC can be one full rate channel (the main stream), which would, in conditions of no packet loss, be the source of the sound reproduction, plus two monaural channels (the sub-streams). In the case a full rate channel packet was dropped, it would be replaced by the two monaural channels, used to reproduce the stereo audio stream, while it would take the loss of both the main rate stream and one of the two sub-streams before the reproduction quality dropped to monaural. This scheme provides the same protection against dropouts as the full channel reproduction SEC, but at a cost of only a 40% increase in bandwidth required, compared to the 200% increase required by the full channel reproduction SEC.

Page 19, delete the second paragraph, page 20 delete the first nine lines, and insert therefor:

A TTP can implement SEC in its initial broadcasts through, for example, transmission of four separate sub-streams :

- A13
- The Main sub-channel (M-channel), a joint normal stereo encoding at 160 kilobits per second (kbps), transmitted MDP / 2 seconds in advance of real time (i.e., the time at which the transmissions are intended to be played).
 - The Immediate sub-channel (I-channel), a mono encoding at 32 kbps, transmitted 1 second in advance of real time.

PRELIMINARY AMENDMENT
U.S. Appl. No. 09/595,013

- A13
- The Delayed sub-channel (D-channel), a mono encoding at 32 kbps, transmitted MDP seconds in advance of real time.
 - Although not required for SEC, in any cases it would be useful to also provide a Text sub-channel (T-channel), transmitting the ASCII text required for the advertising crawl bar and any required control information, transmitted MDP/2 seconds in advance of real time at 5 kbps.
-

Page 20, delete line 17 and insert therefor:

- A14
- G_2 : Lower rate service: the D channel, totaling 32 kbps.
-

Page 21, delete the second paragraph and insert therefor:

A15

When changing channels, the new G_3 group is joined immediately, and playback starts after 3 seconds in the standard PIM-SM. At that time the other groups, if any, in the service are joined.

Page 23, delete the second paragraph and insert therefor:

A16

In normal unicast operation with TCP on an Internet, source based congestion avoidance is in place as the transmission rate is lowered if there is evidence of congestion; namely notification of lost packets by the receivers. This form of congestion avoidance is not possible with large scale multicasts, as congestion events would in general not occur for all receivers simultaneously. These separate groups can be used to implement Receiver-Based Congestion

PRELIMINARY AMENDMENT
U.S. Appln. No. 09/595,013

A16
avoidance (RBCA) based on the amount of time the primary channel has to be replaced by lower rate information. If this occurs for more than RBCA Threshold Ratio proportion of the time in a RBCA threshold interval, the receiver shall implement RBCA. The default values for the RBCA Threshold Ratio is 50% and for RBCA Interval is 5 minutes.

Page 23, delete the third paragraph and insert therefor:

A17
In PIM-SM v.2 a multicast group leave is supposed to be completed in no more than 3 seconds. The MTN receivers, if receiving Group S_i , with $i < 4$, can execute receiver based congestion avoidance by going from S_i to service S_{i+1} by:

Page 24, delete the last paragraph and insert:

A18
While various implementations of the inventive system and model for multicast peering, including Staggered Erasure Correction and Multicast Firewall, have been described in detail, a skilled artisan will readily appreciate that numerous other implementations, particularly those where establishing a multicast transmission is desired, are possible without departing from the spirit of the invention.
